HIGHLY CONFIDENTIAL
FILED UNDER SEAL

UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK

VIACOM INT'L INC., ET AL., Plaintiffs, v. YOUTUBE, INC., ET AL., Defendants	ECF Case Civil No. 07-CV-2103 (LLS)
THE FOOTBALL ASSOCIATION PREMIER LEAGUE LIMITED, ET AL., on behalf of themselves and all others similarly situated, Plaintiffs, v. YOUTUBE, INC., ET AL., Defendants.	ECF Case Civil No. 07-CV-3582 (LLS)

DECLARATION OF DAVID KING IN SUPPORT OF DEFENDANTS' MOTION FOR SUMMARY JUDGMENT

I, DAVID KING, pursuant to 28 U.S.C. § 1746, hereby declare as follows:

1. I work at Google as a Product Manager for YouTube's Content Identification system. I have held the title of Product Manager since I started working at Google in January 2007. Before joining YouTube, I worked for Real Networks, where I was the Senior Director of Content Operations. In that capacity, I ran Real Networks' content-licensing systems and worked closely with record labels and music publishers. I hold a B.A. degree in history from Princeton University.

2. My job responsibilities at Google have focused primarily on managing the development and implementation of YouTube's content-identification technology and its associated rights-management features. Under my supervision, YouTube has invested tens of thousands of person-hours and million of dollars building state-of-the-art video and audio content-identification tools intended to help rights holders better control the usage of their materials on YouTube. In addition to the content-identification tools that have been custom-build by YouTube engineers, YouTube has also licensed audio-identification technology from a company called Audible Magic.

3. These technologies comprise the core of what we refer to as "Content ID." Content ID is a suite of tools that YouTube makes available free of charge to rights holders to make it easier for them to identify videos on YouTube that may contain their content and to instruct YouTube what they want done with those videos. Content ID uses advanced audio- and video-identification technology to scan every new video that users attempt to upload to YouTube—and all videos already posted on the service—against an ever-growing library of reference material supplied by participating copyright holders.

YouTube's Use of Audio Identification Technology

4. When I started working at YouTube in January 2007, YouTube's efforts to use digital technology to assist rights holders had been underway for some time. At that time, I was informed that YouTube had already implemented "MD-5 hash" technology to prevent users from uploading videos that were identical to

videos that had been taken down after a request from a copyright owner. Also, I learned that in October 2006, YouTube had signed an agreement with Audible Magic to license audio identification technology to help rights holders more easily identify their content on YouTube.

5. In early 2007, Audible Magic was among the most established providers of audio-based content-identification technology. Audio-based content identification (sometimes called "audio fingerprinting") is one way of using digital information to try to match unknown files posted to websites such as YouTube to content that may be owned by rights holders. At a high level of generality, an audio fingerprinting technology like Audible Magic works by generating a digital "fingerprint" of the audio track of the unknown file (the "probe") and comparing that probe against an existing database of "reference" files that correspond to content supplied by rights holders. Given the way the technology works, unless a reference file corresponding to a given copyrighted work is in the Audible Magic database, Audible Magic will not be able to match probe files to that work.

6. At the time that YouTube licensed Audible Magic's technology, Audible Magic had strong support from the music industry, particularly major record labels such as Warner Music Group and Universal Music Group. It is my understanding that Audible Magic's technology was developed primarily to help those record labels identify their sound recordings on the Internet. Based on my conversations with Audible Magic, I understood that, as of late 2006 and early 2007, virtually all of the

reference files that Audible Magic was maintaining in its database were from sound recordings owned by major record labels.

7. YouTube used Audible Magic's technology as an important component of our new "Claim Your Content" system ("CYC"). CYC was a predecessor to Content ID; it was a platform that enabled participating rights holders to "claim" videos containing their content that users had uploaded to YouTube. Audible Magic was one way that participating content owners could find videos that they wished to claim. Once it found a video, a rights holder could apply one of three policies that instructed YouTube what to do in response to the match. First, the rights holder could "block" a claimed video, that is, instruct YouTube to remove the video from YouTube. Second, the rights holder could "track" the video, that is, leave it up on YouTube and receive analytics and other information about how it was being viewed. Third, the rights holder could choose to "monetize" the video, that is, leave it up on YouTube and share in revenue associated with advertising that would be displayed on the page where users watched the video.

8. CYC launched in beta form in February 2007. The first copyright holder to use Audible Magic to "claim" a video on YouTube was Universal Music Group on February 14, 2007. In the weeks and months that followed, a number of other rights holders signed up to participate in the CYC program and to use Audible Magic to help identify videos containing their content. All told, between February 2007 and December 2009, approximately 50 different rights holders used Audible

Magic to claim videos on YouTube. YouTube did not charge rights holders to sign up for or to use Audible Magic.

9. Rights holders using Audible Magic on YouTube were free to apply whatever usage policy they wished in the event of a match. YouTube's policy was to make CYC (including Audible Magic) open to all rights holders who wanted to use it, regardless of whether the rights holder was doing so in order to block its content from appearing on YouTube or to claim videos for the purpose of monetization.

10. There were multiple rights holders that used Audible Magic solely to block videos. [redacted] But most rights holders who used CYC chose instead to embrace the promotional opportunities that YouTube provided by allowing the videos they claimed to appear or remain on the service.

YouTube's Development of Video Identification Technology

11. Although the audio-based content-identification technology that Audible Magic provided was useful, particularly in helping the owners of sound recordings identify their content, it had certain limitations in reliably matching against certain kinds of *video-based* content. For example, most television programs and motion pictures include embedded music that is owned by someone other than the entity that owns the TV program or motion picture itself. Particularly because YouTube had entered into carefully negotiated agreements with most of the major record labels to allow their sound recordings to appear on

YouTube, using audio-based content identification to identify television programs and movies was likely to lead to conflicting claims for the same piece of content. For example, the television show "CSI" uses the song "Who Are You" by The Who in its opening credits. An audio-based content identification system will not reliably be able to distinguish a video clip of the opening credits of CSI from a music video of "Who Are You" (or clip from a movie using the same song). Based on our experience with Audible Magic, we found that audio-only matching for video content resulted in confusion and inaccuracy.

12. In addition, many audio-visual works have a variable soundtrack, which can minimize the utility of audio-based content identification. Sporting events, for example, are often broadcast in different languages and with different commentators. And the background "stadium sounds" for many different sporting events (things like crowd noise and whistles, for example) are often quite similar and difficult to distinguish from each other. Audio-based content identification technology therefore would often be unreliable for identifying such works.

13. For these reasons, it was my belief (and the belief expressed to me by others on my team) that the most effective and reliable content identification technology for a video website like ours would be *video-based* content identification (sometimes called "video fingerprinting"). Video-based content identification works much like audio fingerprinting, with the important difference that the former uses the *video* channel of the probe file in identifying potential matches. By looking at the video channel, rather than just the audio channel, video-based content

identification solves some of the key problems with using audio fingerprinting to try to identify audio-visual content such as television shows and movies.

14. From the very start of my tenure at Google, I saw video-based content identification as a meaningful new way for YouTube to further help rights holders find videos on YouTube that might contain their content. I also viewed the development of video-based content identification as an exciting technical challenge that would lead the way for other user-generated content websites, none of which had implemented such technology. Accordingly, in January 2007, almost immediately after I began working at Google, I made the decision that YouTube should build a video-based content identification tool to supplement (and ultimately supplant) our use of Audible Magic. It was expressed to me that the decision to build that technology—which we came to call "Video ID"—had the full support and encouragement of Google and YouTube management, including Eric Schmidt (the CEO of Google) and Chad Hurley (the CEO of YouTube).

15. Although the Video ID project formally began in January 2007, Google's work on video-based content identification technology started well before I joined YouTube. When I first began investigating the feasibility of YouTube developing its own video-based content-identification tool, I learned that YouTube engineers had already been in active discussions with Google engineers about groundbreaking video-recognition technology that Google had been working on for several years. I was told to pick up those discussions as part of my new job responsibilities. Based on those conversations, my team and I were convinced that

the work that Google had already done could serve as the backbone for a state-of-the-art video-based content identification system that could be built specifically for YouTube.

16. There were several reasons why YouTube decided to develop its own content-identification technology, as opposed to relying on technology from an external vendor. *First*, at that time there was no commercially available video-based content identification technology for use on websites like YouTube. Although there were a few companies that were testing early versions of such technology, none of them had a product that had actually been commercially deployed on any website. Nor were we confident that any of these third-party vendors was or would soon be in a position to offer video-based content identification technology that could reliably and efficiently operate on a site that handles the volume of video uploads to YouTube. As of early 2007, YouTube's scale of operations dwarfed that of any other video website, and that scale posed a significant technical and operational challenge to any content identification system. There was nothing available on the market, or even on the horizon, that seemed up to that challenge. *Second*, as I mentioned above, before its acquisition of YouTube, Google had already done significant work on video-identification technology, which we believed could be adapted to YouTube's needs without requiring us to build a product from scratch. That led us to believe that we could develop our own video-based content identification system more quickly and effectively than could any third party. *Third*, by building the technology ourselves, we could design it specifically to run on YouTube's systems.

That would ensure that Video ID was efficiently integrated into YouTube's architecture, which we thought would make it faster, more reliable, better able to operate at scale, easier for rights holders to use, more adaptable to their needs, and ultimately more cost-effective, than any third-party technology.

17. I was responsible for leading the team that built and implemented Video ID. Full-scale development efforts began in January 2007. By July of that year, we had a product that was ready for initial testing by content owners. We were able to make such rapid progress in large part because of the skilled and dedicated team of engineers who were devoted to the project. Between January and October 2007, we had between 15 and 20 engineers and other technical personnel working full or part time on the project, not to mention dozens of other people involved in business, legal, and operational issues relating to the development and implementation of Video ID. All told, I would estimate that YouTube devoted more than 50,000 person-hours to building Video ID and preparing it for its public launch.

18. YouTube invited a group of major content owners to participate in pre-launch ("beta") testing of Video ID. Those tests occurred between July and September 2007. After getting feedback from the content owners who participated in those beta tests (including Viacom) and making further refinements to the technology, YouTube officially launched Video ID in October 2007.

19. Given the nature of my work at YouTube, I pay attention to what other similar websites are doing in the field of content-identification. To my knowledge,

Video ID was the first video-based content identification technology to be deployed on any website dedicated to user-submitted content. YouTube was certainly the first (and I believe remains the only) such website to have developed and launched its own video-based content identification system.

20. In April 2008, YouTube supplemented Video ID by launching Audio ID. Like Audible Magic, Audio ID is an audio-based content identification technology, but because it was custom-built by Google and YouTube engineers specifically for use on YouTube in conjunction with Video ID, it is faster and more efficient in identifying audio content in the YouTube environment. Today, Audio ID and Video ID work together to make up the technological backbone of YouTube's Content ID system. Together, they create a cohesive suite of content-identification technologies that we can confidently make available to identify a wide variety of content on behalf of rights holders around the world.

YouTube's Implementation of Content ID

21. At the time that Content ID launched in October 2007, a wide range of rights holders signed up to take advantage of the tool. Since then, more and more content owners have started using Content ID. We currently have over 1000 content owners worldwide using Content ID to identify their content on YouTube. That includes every major U.S. television broadcaster, movie studio, and record label, as well as most of the largest music publishers and many of the major sports leagues in the United States and abroad. A true and correct list of the content owners signed up to use Content ID is attached hereto as Exhibit 1.

22. YouTube makes (and has always made) Content ID available to rights holders free of charge. We have also worked hard to make the tool easy to sign up for and use. The only requirement to get access to Content ID is to sign a short agreement that ensures that YouTube and the participating rights holder have a mutual understanding of each party's rights and responsibilities in using this powerful technology for the purpose that it was intended: to assist rights holders with identifying and managing their own content, and not laying claim to content they do not own.

23. Content ID works by identifying videos on YouTube that match reference files supplied by participating rights holders. The library of reference files that YouTube currently maintains for that purpose is extensive. As of December 2009, it consisted of approximately 3 million reference files provided by participating rights holders.

24. If Content ID identifies a video as matching one of those reference files, the rights holder has several options for what it can instruct YouTube to do with that video: (a) block/remove the video; (b) allow the video to appear and share any revenue generated from advertising shown alongside it; (c) allow the video to appear, but with no monetization. Rights holders can designate their preferred policies in advance, which are then automatically applied by the Content ID system.

25. Rights holders can also designate different policies to apply to different regions, for example, instructing YouTube to block a given video from being shown to users in the United States, but allowing it to be displayed to users in Europe.

Rights holders can also condition their policy choices based on other factors, such as the duration or proportionality of the match. For instance, a rights holder could instruct that videos that match more than 5 minutes of one of their reference files be blocked, while shorter matches be monetized. That is consistent with the basic ethos of Content ID, which is to give maximum flexibility and choice to rights holders about whether and how they want their content to appear on YouTube.

26. Since its launch in October 2007, every video that anyone attempts to post on YouTube has been screened using Content ID. Given the volume of new video uploads to YouTube, Content ID scans an enormous volume of video. As of December 2009, Content ID was scanning approximately 20 *hours*-worth of new video each *minute*. For each new video uploaded, Content ID generates a result (match or no-match) within a matter of seconds. Additionally, at the request of certain rights holders, YouTube has configured Content ID to scan new videos before they are actually posted. Thus, if a user attempts to upload a video that results in a match and the relevant rights holder has set a policy of "block," the system is designed to prevent that video from ever making its way onto YouTube.

27. In addition to doing "pre-publication" scanning, Content ID also scans the full back catalogue of videos that are already posted on YouTube (currently, more than 200 million videos). When combined with the daily new uploads, that means that Content ID scans approximately 100 *years*-worth of video each day.

28. YouTube currently has a team of approximately 40 technical staff (engineers and product managers) working on Content ID. Since Content ID first

launched, YouTube has continued to improve the product, working to make the technology even faster, more reliable, and more scalable. The technology that supports the Content ID platform is robust. Video and Audio ID are able to identify videos as matching a content owner's reference material even where the video may be significantly distorted or altered from the original file.

Viacom's Use of Content ID

29. Viacom is an active user of Content ID. Viacom was one of the initial group of rights holders who participated in the "beta" testing of Video ID before its public launch. After completing that testing, Viacom signed up to use Content ID in February 2008.

30. Since it began using Content ID, Viacom has provided thousands of reference files to YouTube for content that Viacom wished to be matched using the tool. It is my understanding that Viacom has not supplied YouTube with reference files corresponding to all of the content that Viacom owns. For example, Paramount has provided reference files for only 50 or so of its motion picture titles. Of course, any works that Viacom does not provide to YouTube as a reference file cannot and will not be matched using Content ID.

31. For the majority of the reference files it has provided for use in the Content ID system, Viacom has instructed that YouTube apply a policy of "block" to matching videos. Consistent with those instructions, Content ID has blocked or taken down videos identified as matching those Viacom references. For some of its reference files, however, Viacom has instructed YouTube to apply a policy of "track."

Based on Viacom's instructions, Content ID has allowed videos identified as matching those Viacom references to be posted to or remain on YouTube and provided information to Viacom about how YouTube users are engaging with the matching videos.

I certify under penalty of perjury that the foregoing is true and correct.

Dated: San Bruno, California
March 1, 2010

David King